arXiv:1612.03777v1 [cs.CV] 12 Dec 2016

Next-Flow: Hybrid Multi-Tasking with Next-Frame Prediction to Boost Optical-Flow Estimation in the Wild

Nima Sedaghat
University of Freiburg
Germany
nima@cs.uni-freiburg.de

Abstract

CNN-based optical flow estimators have attracted attentions recently, mainly due to their impressive speed. As successful as they've been on synthetic datasets, they are still far behind the classical methods in real-world scenarios, mainly due to lack of flow ground-truth. In the current work, we seek to boost CNN-based flow estimation in real scenes with the help of the freely available self-supervised task of next-frame prediction. To this end we train the network in a hybrid way, providing it with a mixture of synthetic and real videos. With the help of a novel time-variant multi-tasking architecture, the network decides which of the tasks to learn at each point of time, depending on the availability of ground-truth. Our proposed method improves optical flow estimation in real scenes dramatically. We also experiment with prediction of "next-flow" instead of estimation of the current flow, which is intuitively more related to the task of next-frame prediction and improve the results even more.

We report and evaluate results both qualitatively and quantitatively. The latter is done by training a single-stream action classifier on the estimated flow fields on UCF101 & HMDB51 and demonstrate high improvements of accuracy over the baseline: 10.2% and 9.6% respectively. Also as a side product of this work, we report significant improvements over state-of-the-art results in the task of next-frame prediction.

1. Introduction

Optical flow estimation is one of the key topics of research at the core of computer vision. Since the work of Horn & Schunk [6], most of the work in this area has been based on variational methods and the leading works such as DeepFlow [22] and EpicFlow [17] are all variants of it. But very recently, CNN-based flow estimation methods, have become the center of attention. The works of [19] & [5], later followed by [20] were the first ones to train a ConvNet in a simple end-to-end fashion for optical flow estimation. The impressive advantage of the network-based methods is their inherent speed at the test time, which makes them eligible to be called "real-time" methods.

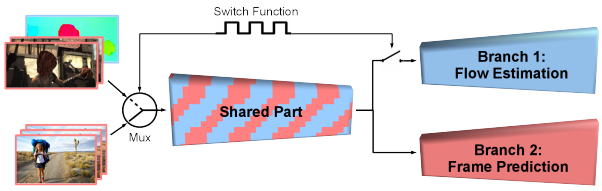

Figure 1: We improve CNN-based optical flow estimation in real videos, by adding the extra *self-supervised* task of future prediction, and training the network with a mixture of synthetic and real-world videos. This combination is done by putting a "multiplexer" at the entry of the network which mixes data from the two sources on a timely basis, at batch level. To make this "hybrid" data digestable by the network, we implement an on-off switch to deactivate part of the network at certain times, and keep this switch in sync with the multiplexer.

Although the network-based methods are relatively successful on synthetic datasets, they are still behind the leading classical methods in real world by a large margin, mainly due to lack of training data with optical flow groud-truth on real videos. To overcome this issue, many methods focus on creation of better synthetic datasets [5, 14]. In contrast to such solutions, we seek to deal with the problem by making the network "see" the real data during training. But this is not a straight-forward task. [20] applies a variational method on real data to generate "semi-truth" data, while [1, 24] approach the problem by reformulating the task as an unsupervised one for the network. The latter works both use cost functions based on the constant intensity assumption, during training.

We also feed the real data to the network in an unsuper-

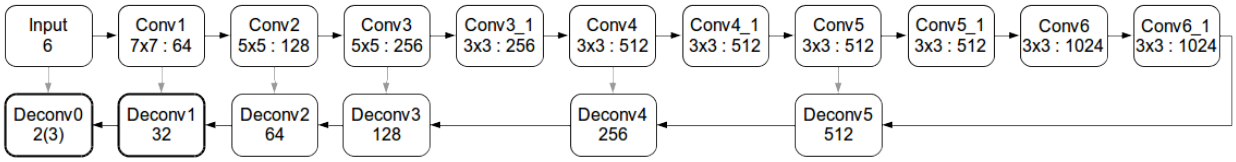

Figure 2: Template architecture used for the flow estimator and frame predictor networks. For the sake of brevity, the small refinement branches introduced in [5] are not diplayed here. The flow estimation and frame prediction networks differ in the number of outputs of Deconv0 – and the hidden lower-resolution outputs. The deconvolutional layers are designed such that the output of each layer is of the same size (resolution) as the corresponding convolutional layer.

vised manner. However we use the self-supervised task of next-frame prediction as an auxiliary task for the network, to be able to leverage the real data without the need for flow annotations. Video prediction has been widely experimented with recently [13, 15, 7, 4, 11, 23, 12]. Although some of these works, such as [13], focus on prediction as the main objective, most of the others use it as an auxiliary task. [15] learns optical flow by trying to warp the current frame to the next one.

We propose a hybrid multi-tasking scheme, which makes use of a mixture of two different types of data: synthetic video with flow ground-truth vs. real video without. Only very recently, and concurrently with our work, [10] proposes multi-tasking with diverse datasets.

Contributions Our contributions in this work can be summarized as:

- Implementation of a multiplexer layer which together with an on-off switch inside the loss layer, form a time-variant architecture and make hybrid training of the network on datasets of different type possible.
- Simultanous training on supervised and self-supervised tasks to highly improve CNN-based optical flow estimation in real scenes.
- Prediction of the future flow instead of current flow as a more compatible task with video prediction.
- State-of-the-art next-frame prediction results as a single task, just using L1 loss.

2. Methods

In sections 2.1 and 2.2 we first introduce our flow estimation and next-frame prediction components as individual networks. Then in section 2.3 we show how these architectures are combined together with the time-dependent switching mechansim, to form the proposed hybrid multi-tasking architecture. Finally in 2.4 we show how "current flow" estimation is replaced by a "future flow" prediction.

2.1. Flow Estimation

The flow estimator network we use, has a fully convolutional encoder-decoder architecture, which builds on the one introduced in [5]. We choose to stick to this particular kind of architecture to be able to clearly analyze the effects of the added auxiliary task and data type. However we extend the architecture of [5], by adding two more deconvolutional layers, to get the full-resolution flow-field directly from the network. We do this to keep this architecture in harmony with the frame predictor. We use the endpoint error loss (EPE) for training of this component.

The architecture illustrated in Figure 2 details the contraction and expansion in the encoder and decoder components respectively. In Figure 2 we omit the extra details regarding the so-called "refinement" steps (Figure 3 in [5]), to focus on the main contributions of the current study. However we depict the skip connections between the Conv/Deconv pairs, to show how they are extended to the two new Deconv layers – note the link from Input to Deconv0. Also note that in Figure 2 number of outputs of Deconv0 is set to 2 for the task of flow estimation.

2.2. Next-Frame Prediction

We adapt the architecture introduced in the previous section for flow estimation, to the task of next-frame prediction. To this end, we simply change the number of outputs of Deconv0 to 3, to represent 3 channels of an RGB image – instead of a 2-channel flow frame. As suggested in preivious work such as [13], we avoid using an L2 loss for this task to avoid blurriness of the results, and use L1 loss for training. As reported in section 3, we experiment with different number of frames as input. However in the multi-tasking scheme, we only use a 2-frame set-up to be compatible with the paired task of flow estimation.

2.3. Mixing It Up

Multi-Tasking If we were to train just on synthetic data, where flow ground-truth is available, a simple multi-tasking approach would be as easy as designing a network to produce two outputs in parallel, one for each task. The two tasks would share some parameters and during training we would minimize the total loss. We basically follow the same scheme for combination of the two tasks, such that the two networks share the encoder (convolutional) layers and are independent in their decoding (deconvolutional) branches (Figure 1).

In the problem at hand, however, there are two complexities to overcome. First, the data is coming from two differet sources and we should find a way to mix them. Secondly, unlike synthetic data, the real data does not come with flow ground truth, and the question is what the flow-related part of the network should be doing when it faces real data.

Hybrid Data For mixing the data, we implemented a time-based multiplexer layer whose task is to alternate between the two data "channels" (Figure 1). The multiplexer does this alternation at the batch level: data in a single batch is guaranteed to be taken from a single source. We can formulate the job of the multiplexer as:

$$\mathcal{B}^{(i)} = (1 - s^{(i)})\mathcal{B}_1^{(i)} + s^{(i)}\mathcal{B}_2^{(i)} \tag{1}$$

where $\mathcal{B}_1^{(i)}$ & $\mathcal{B}_2^{(i)}$ are the data blobs from datasources 1 & 2 at the i^{th} iteration, respectively, and $\mathcal{B}^{(i)}$ is the output of this layer which is fed to the network.

$s^{(.)}$ is the *switch* function as follows:

$$s^{(i)} = \lfloor \frac{i \mod (n_1 + n_2)}{n_1} \rfloor \tag{2}$$

in which, n_1 and n_2 define the number of cycles dedicated to each datasource. The output is a periodic signal consisting of n_1 0's and n_2 1's respectively.

Time-Variant Architecture As for the second issue, we implement an on-off switch whose job is to deactivate the flow-related part of the network when there's no flow annotation, so that no learning happens during those iterations (Figure 1). The On-Off switch is in practice implemented inside the loss layer such that during the off state, both loss and loss gradient are set to zero.

The total loss at the i^{th} iteration is computed according to:

$$\mathcal{L}^{(i)} = w_1\mathcal{L}_1^{(i)}s^{(i)} + w_2\mathcal{L}_2^{(i)} \tag{3}$$

in which $\mathcal{L}_1^{(i)}$ & $\mathcal{L}_2^{(i)}$ are the Flow and Frame estimation respectively (1), with their assigned weights, w_1 & w_2.

We keep this switch in sync with the multiplexer layer, to achieve the desired functionality: the network learns both tasks during the "synthetic iterations", but postpones updating the flow-only branch when there's no ground-truth. Of course the shared part of the network remains activated all the time.

We set the loss weights such that w_1/w_2 is equal to σ_2/σ_1 where σ's are the estimates of the variances of input data (frame vs. flow) and are computed over a subset of the training sets.

In our main experiments we dedicate an equal cycle of 1 to both synthetic and real data sources, but we also provide an anaysis on the effect of different cycle ratios on the quality of the output flow field.

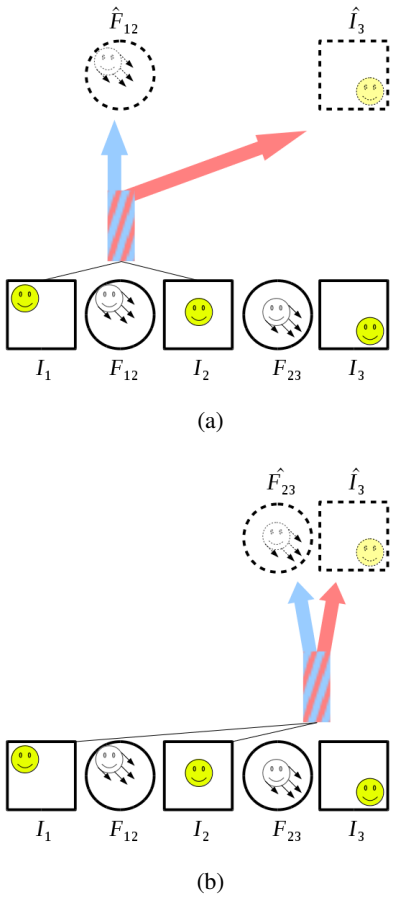

Figure 3: Comparison of the two introduced multi-tasking schemes. a) illutstrates combination of Flow Estimation with Next-Frame prediction whereas in b) the Next-Flow Prediction replaces the task of current flow estimation. Each I_k represent a single frame in a video sequence of length 3 and F denotes a flow field. In both scenarios only I_1 and I_2 are the inputs to the network. Therefore the terms "current flow" and "next-flow" refer to F_{12} & F_{23} respectively.

2.4. Next-Flow Prediction

In a multi-tasking scheme, for the combination to yield significant improvements in the results, the two tasks need to be "related" [3].

In context of the current work, we hypothesize that prediction of the "next-flow" (i.e. the future flow to come), may have more in common with the task of next-frame prediction. Figure 3 compares the two combinations, and gives an intuition on how the two "prediction" tasks might be in more harmony.

More formally: The two tasks share the encoder component of the network, that maps $(I_1, I_2) \mapsto z$, where z is the internal representation to be learned by the network. This mapping is affected by both tasks during the backward pass and for multi-tasking to make sense, we expect that some part of the features learned by the encoder be beneficial to both of the target tasks. We hypothesize that the pair of $(z \mapsto I_3, z \mapsto F_{23})$ have more to share, compared to $(z \mapsto I_3, z \mapsto F_{12})$ not only due to both being "prediction" tasks, but also because in obtaining the future flow, the network needs to learn, at least implicitly, about the future frame. This is not the case for the current flow (Figure 3). This hypothesis is supported by experimental results in the following.

3. Experiments & Results

We evaluate the optical flow estimation capabilities of the proposed networks both qualitatively and quantitatively, where possible.

3.1. Datasets

For hybrid training of the network we need 2 datasets. We use the the so called "FlyingThings3D" dataset of [14] as the synthetic data source. It consists of more than 20K training images and is a more suitable choice for training a ConvNet from scratch, compared to the famous Sintel dataset with only 1064 training samples [2]. For the real-world source of data, we use a subset of the Sports1M dataset [9]. The subset includes all the videos with a file size less than 5 MBytes, which amounts to more than 220K videos and 220M frames. We make the selection list available online.

We use the UCF101 [18] and HMDB51 [8] datasets for evaluation of the optical flow on real datasets. These are both rather big datasets containing more than 2M & 600K frames respectively. Also to analyze the effect of the proposed method on synthetic data, we use the public part of the Sintel dataset [2] along with the test subset of FlyingThings3D [14].

To compare the performance of our next-frame predictor to published work, we use a subset of UCF101 [18] consisting of 387 videos, as suggested by [13].

	EPE	
	Flying Things 3D (test)	Sintel (train)
FlowNet	2.4	4.7
FlowNet+NextFrame	2.3	4.8

Table 1: End point error (EPE) measured on two *synthetic* datasets before and after multi-tasking. As expected the values are not changed significantly.

3.2. Direct Analysis of The Flow Fields

In Figure 4 we visualize some of the flow fields estimated using our method, and compare to the baseline results and the (pseudo-)ground-truth. On the real-world scenes, our flow fields clearly outperform the baseline in preserving shapes of the objects. We believe this sharpness is a result of asking for pixel-level accurate results in the auxiliary task of frame prediction.

On synthetic datasets, we do not expect the results to get any better, compared to the baseline network which is fully trained on scenes of the same type – this is confirmed by the results in table 1. However, there are still many synthetic samples on which our network has made noticable improvements. Figure 4 depicts some of them at the bottom.

3.3. Indirect Evaluation – Action Classification

As real-world videos do not come with optical flow ground-truth, it is not possible to directly quantify the optical flow estimation performance in large scale. Thus we use the task of flow-based action classification, as an indirect quantitative evaluator. We use the action classifier network of [21] and train/test it with different flow estimators. Table 2 contains the results of this evaluation. To analyze the necessity of the multi-tasking scheme It is clear that just pre-training with the second task/ We improve the baseline results by a large margin (9.7% on UCF and 8.9% on HMDB), by just adding the second task and datatype ("FlowNet+NextFrame Multi-Tasking"). We achieved even more improvements by replacing current Flow with NextFlow, which confirms the hypothesis of section 2.4.

We also tried a pre-training/fine-tuning scenario in which the network is initially trained for the frame prediction task (on real data), and then fine-tuned with the main task ("FlowNet pre-trained with NextFrame"). Results confirm that the multi-tasking scheme has been necessary to get the most out of the auxiliary task and real data.

Results of two of the state-of-the-art classical methods, TV-L1 [16] and EpicFlow [17], are still better than the CNN-based results.

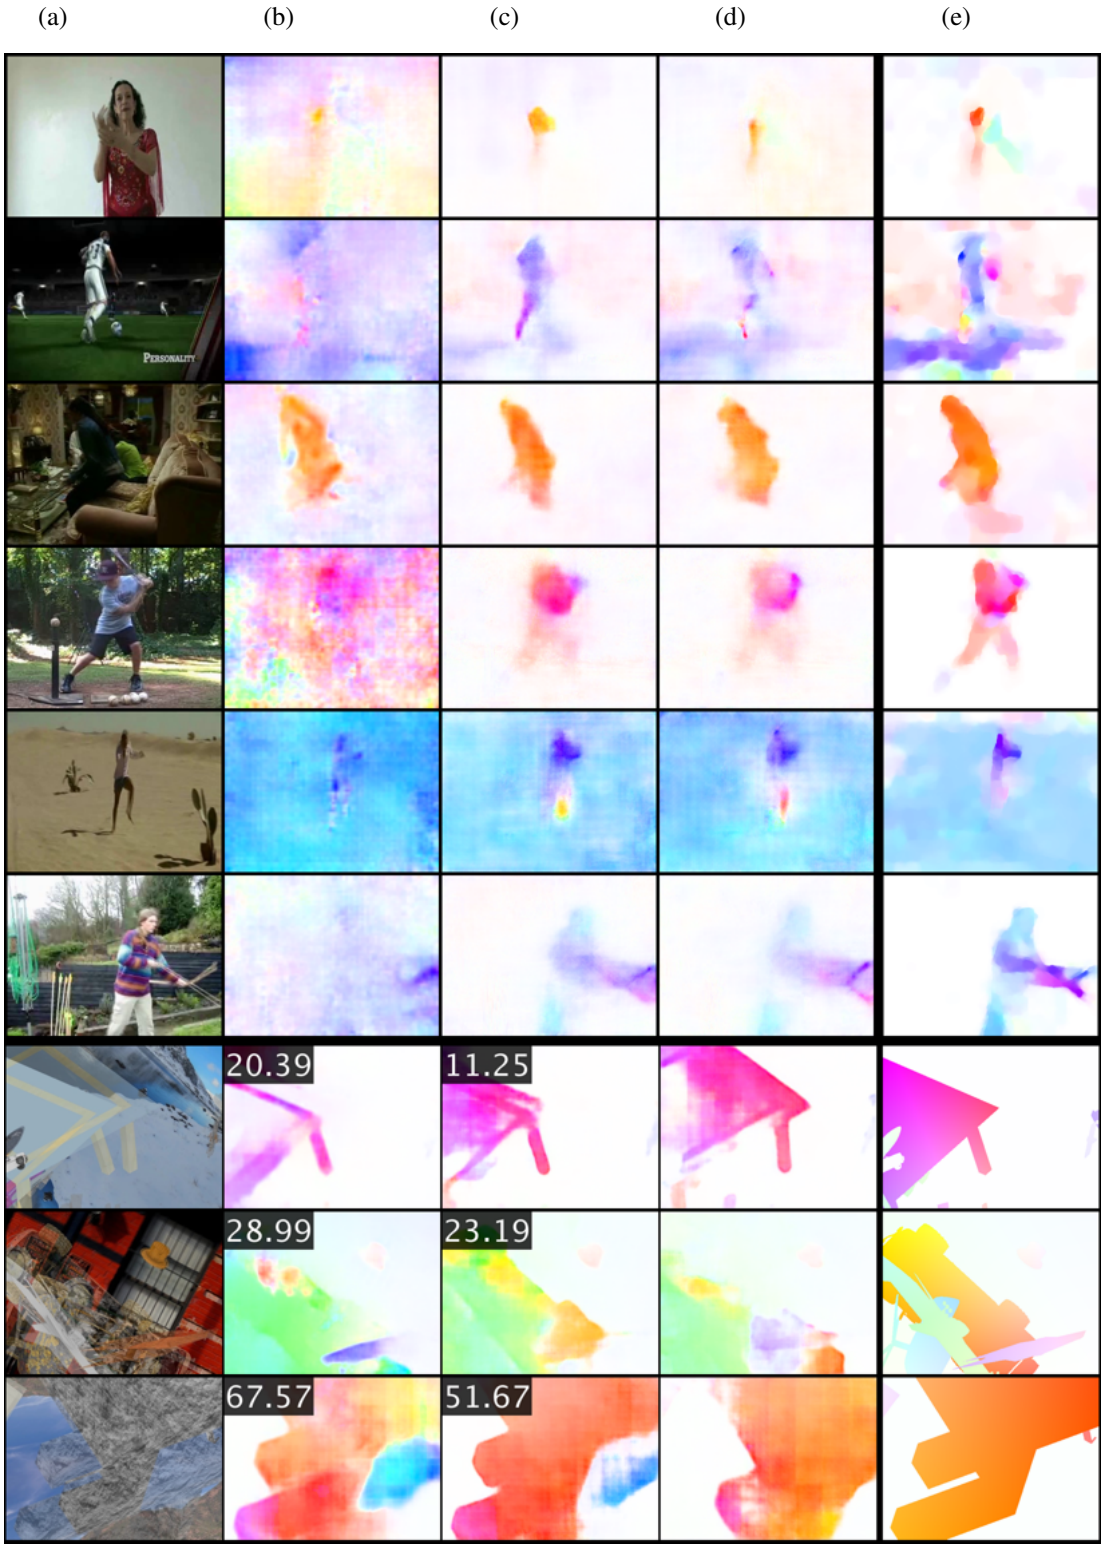

Figure 4: Some samples from the estimated optical flow fields. The three bottom rows are from the synthetic scenes and the rest from real data. a) Overlayed input images b) FlowNet c) FlowNet+NextFrame d) NextFlow+NextFrame e) EpicFlow for real scenes and ground truth for synthetic ones

		UCF101 [18]	HMDB51 [8]
Classical	EpicFlow [17]	82.8%	56.1%
	TV-L1 [16] (as reported in [21])	87.2%	-
CNN based	FlowNet Baseline (as published in [5])	62.0%	38.6%
	FlowNet pre-trained with NextFrame	63.4%	38.4%
	FlowNet+NextFrame Multi-tasking	71.7%	47.5%
	NextFlow+NextFrame Multi-tasking	**72.2%**	**48.2%**

Table 2: Action classification accuracies. Each row contains results of training and testing the action classifier on optical flow generated by a specific method.

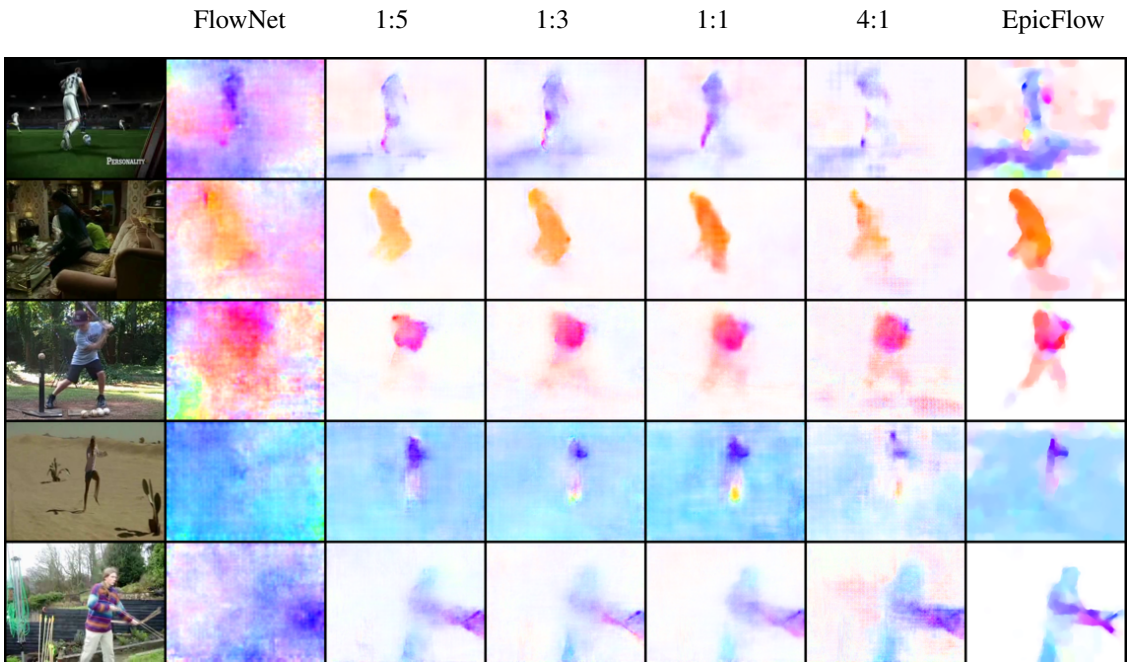

Figure 5: Visualization of the effect of different datasource combination cycles. On top of each column the $^{n_1}/_{n_2}$ ratio is displayed. When there are is no frame prediction involved (FlowNet), the network fails to estimate an acceptable flow field in real scenes. On the other hand if we spend too many cycles on frame prediction, as in the 4:1 column, the network starts forgetting about its main job. We can see the best results in 1:3 or 1:1 columns depending on the sample. Determination of the exact best combination setting would be a hyper-parameter optimization problem.

	EPE	
$^{n_1}/_{n_2}$	Flying Things 3D (test)	Sintel (train)
FlowNet	2.4	4.7
1:5	2.1	4.5
1:3	**2.1**	**4.4**
1:1	2.3	4.8
4:1	3.0	5.7

Table 3: EPE on the two synthetic test datasets for networks trained with different task combination cycles.

	Accuracy (%)	
$^{n_1}/_{n_2}$	HMDB51	UCF101
FlowNet	38.6	62.0
1:5	44.6	70.8
1:3	**49.1**	71.5
1:1	47.5	**71.7**
4:1	44.6	67.0

Table 4: Action classification accuracies for various task combination settings.

3.4. The Impact of Task Combination Cycles on Results

The multiplexer introduced in section 2.3 mixes the two data-sources, assigning n_1 and n_2 cycles (iterations) to each

	Method	Whole Image		Moving Regions	
		Similarity PSNR(dB)	Sharpness (dB)	Similarity PSNR(dB)	Sharpness (dB)
Mathieu *et al*. [13]	L1	22.3	18.5	28.7	24.8
	GDL + L1	23.9	18.7	29.9	25
	Adv + GDL + L1	29.6	20.3	32	25.4
Ours (2-frame)	L1	29.9	20.6	31.9	25.4
Ours (4-frame)	L1	**30.8**	**20.8**	31.9	25.4

Table 5: Next frame prediction on UCF101 [18].

(Equations 1,2). Figure 5 and tables 3,4 represent the results with various $^{n_1}/_{n_2}$ (or real/synthetic iterations). It is seen that we may get slightly better results by increasing the number of flow learning cycles, e.g. from 1:1 to 1:3. However by going higher than this ratio, the effect of the auxiliary task starts to vanish, and the flow results are degraded.

It is also interesting to see that hybrid training can, in some settings, significantly help for synthetic data too.

3.5. Next-Frame Prediction as a Single Task

We also train and test our Next-Frame predictor as an independent single-tasked network. Table 5 shows a comparison of our method with [13] – which to the best of our knowledge is the current state-of-the-art in next-frame prediction on UCF. We believe that the ability to predict the static regions of the image is as important as prediction on the dynamic areas, and thus our results include evaluations on the whole image as well as only on the moving pixels.

Wihtout using auxiliary cost functions for the sake of sharpness, and with a single L1 loss, we obtain equivalent results on the moving regions of the image, and significantly better results on the whole image. This means that the network is doing a better job 'applying the motion" only to the dynamic areas and keeping the static areas intact.

4. Conclusions

Needless of looking for more synthetic data, we have boosted the performance of a CNN-based optical flow estimator using the freely available data from real-world videos. The quality of the estimated optical flow on real-world scenarios is dramatically improved, while on the synthetic data it remains more or less the same. There's still a noticable distance to the performance of classical flow estimation methods, which leaves room for more work on CNN-based methods. Our suggested scheme can be plugged into future network-based flow estimators which may use more sophisticated architectures or even benefit from more synthetic data, with little to no modifications. As another possible line of work, one may incorporate auxiliary tasks other than future prediction.

References

[1] A. Ahmadi and I. Patras. Unsupervised convolutional neural networks for motion estimation. *arXiv:1601.06087 [cs]*, Jan. 2016. 1

[2] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, Oct. 2012. 4

[3] R. Caruana. Multitask Learning. In S. Thrun and L. Pratt, editors, *Learning to Learn*, pages 95–133. Springer US, 1998. 4

[4] C. Finn, I. Goodfellow, and S. Levine. Unsupervised Learning for Physical Interaction through Video Prediction. *arXiv:1605.07157 [cs]*, May 2016. 2

[5] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning Optical Flow with Convolutional Networks. *arXiv:1504.06852 [cs]*, Apr. 2015. 1, 2, 6

[6] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1):185–203, Aug. 1981. 1

[7] D. Jayaraman and K. Grauman. Look-ahead before you leap: End-to-end active recognition by forecasting the effect of motion. *arXiv:1605.00164 [cs]*, Apr. 2016. 2

[8] H. Jhuang, H. Garrote, E. Poggio, T. Serre, and T. Hmdb. A large video database for human motion recognition. In *Proc. of IEEE International Conference on Computer Vision*, 2011. 4, 6

[9] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale Video Classification with Convolutional Neural Networks. In *CVPR*, 2014. 4

[10] I. Kokkinos. UberNet: Training a 'Universal' Convolutional Neural Network for Low-, Mid-, and High-Level Vision using Diverse Datasets and Limited Memory. *arXiv:1609.02132 [cs]*, Sept. 2016. 2

[11] W. Lotter, G. Kreiman, and D. Cox. Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning. *arXiv:1605.08104 [cs, q-bio]*, May 2016. 2

[12] R. Mahjourian, M. Wicke, and A. Angelova. Geometry-Based Next Frame Prediction from Monocular Video. *arXiv:1609.06377 [cs]*, Sept. 2016. 2

[13] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv:1511.05440 [cs, stat]*, Nov. 2015. 2, 4, 7

[14] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. *arXiv:1512.02134 [cs, stat]*, Dec. 2015. 1, 4

[15] V. Patraucean, A. Handa, and R. Cipolla. Spatio-temporal video autoencoder with differentiable memory. *arXiv:1511.06309 [cs]*, Nov. 2015. 2

[16] J. S. Pérez, E. Meinhardt-Llopis, and G. Facciolo. TV-L1 optical flow estimation. *Image Processing On Line*, 2013:137–150, 2013. 4, 6

[17] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow. pages 1164–1172, 2015. 1, 4, 6

[18] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv:1212.0402 [cs]*, Dec. 2012. 4, 6, 7

[19] D. Teney and M. Hebert. Learning to Extract Motion from Videos in Convolutional Neural Networks. *arXiv:1601.07532 [cs]*, Jan. 2016. 1

[20] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Deep End2End Voxel2Voxel Prediction. *arXiv:1511.06681 [cs]*, Nov. 2015. 1

[21] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision*, pages 20–36. Springer, 2016. 4, 6

[22] P. Weinzaepfel, J. Revaud, Z. Harchaoui, and C. Schmid. DeepFlow: Large Displacement Optical Flow with Deep Matching. pages 1385–1392, 2013. 1

[23] T. Xue, J. Wu, K. L. Bouman, and W. T. Freeman. Visual Dynamics: Probabilistic Future Frame Synthesis via Cross Convolutional Networks. *arXiv:1607.02586 [cs]*, July 2016. 2

[24] J. J. Yu, A. W. Harley, and K. G. Derpanis. Back to Basics: Unsupervised Learning of Optical Flow via Brightness Constancy and Motion Smoothness. *arXiv:1608.05842 [cs]*, Aug. 2016. 1